

MONITORING TRAFFIC IN PACKET NETWORKS USING THE SLIDING WINDOW PROCEDURE WITH SUBWINDOWS

TECHNICAL FIELD

The present invention relates to monitoring traffic in packet networks, in particular to policing and especially metering functions.

BACKGROUND

In an IP-network information is transmitted between hosts through routers placed at the edge of and inside the network, see Fig. 1. An IP-network is an example of a packet network which generally uses a multitude of switching elements connecting terminals placed at the edge of network. A packet based network can generally support variable bit rate services, such as different codings of speech and data services, much better than an STM (Synchronous Transfer Mode) network, since statistical multiplexing can be used. The multiplexing gain in packet networks improves e.g. the number of speech/data channels that can be used to transmit information between the terminals. Operation and maintenance of packets networks can be simplified compared to STM networks and the traffic associated therewith be multiplexed with other traffic.

However, when the traffic flow on a link between e.g. two routers in an IP-network is too high, congestion can occur and then some packets must be discarded, since otherwise the real-time demands of some of the traffic will not be fulfilled. The new technologies that allows the use of IP-network also for voice traffic introduce different levels of service for different traffic types. Voice traffic is very sensitive to delay and jitter, but data is more sensitive to loss of packets. Thus, a separation of the levels of QoS (Quality of Service) is needed between voice and data traffic. QoS separation is the capability of a network to provide a better service for selected network traffic, i.e to treat different classes of traffic differently. An architecture providing QoS separation is the Differentiated Services, DiffServ, as proposed by the IETF, the Internet Engineering Task Force. Still, this architecture lacks some performance so there is a need for improved solutions.

One of the main causes of congestion in packet switched networks comprises that the traffic is often bursty. If the traffic was uniform, such as e.g. in circuit switched telephony, it would be easy to predict its behaviour and dimension the networks accordingly. In the packet data case, dimensioning is much more difficult and requires weighting between packet loss, induced delay and efficient usage of bandwidth.

To use the bandwidth efficiently, the bursts must be smoothed out in time. This is done by using buffers in the routers in which packets will be queued until the link is ready - a delay is thus induced in the queue - and it is regarded as implicit traffic shaping. This works well when the average traffic flow over a period of time is less then the egress bandwidth but when the ingress traffic over a period of time exceeds this limit, some packets must be discarded.

Policing is a general term for the process of preventing a traffic stream from seizing more than its share of the resources, i.e. to protect the network from non-trusted users.

Policing can be done for the aggregated flow, for the separate service classes of DiffServ or for the packet flow of individual users. Policing acts in a short time perspective and actions take place locally, at the router level. The policing function consists of two different parts, metering and discarding. Metering is the function that monitors the arrival times
5 of packets of a traffic stream and determines the level of conformance of each packet to a pre-established traffic profile.

The meter function should be able to handle bursty traffic and at the same time utilize the bandwidth on the egress link efficiently. When congestion is indicated, the discarding function is activated, deciding whether to transmit or discard the arriving
10 packets. This decision will depend on several parameters such as aggregated flow, class based flow, QoS class of the packet, etc. Discarding can also take place in a scheduler, i.e. that part of a router which actually forwards the packets, but it is preferable that all of the dropping takes place in the policer because of the unpredictable manner in which packets are dropped in a scheduler.

15 When the traffic to a large extent consists of voice packets, the delay demands are so important that the scheduler should be a priority scheduler, provided that the voice traffic is assigned a high priority traffic class. A priority scheduler always serves the non-empty queue of highest priority. This might cause starvation of the lower priority traffic but it can also imply too long delays to the high priority traffic if too much traffic is let through,
20 i.e. it will cause too long queues in the scheduler.

One known solution to this problem, i.e. how to police the traffic flow in order to fulfil the demands of the different traffic classes, is to use a Token Bucket procedure as the policer, see A. Tannenbaum, "Computer Networks - Third Edition" Prentice Hall 1996, pp. 381 - 384. The Token Bucket procedure monitors the packets of the incoming
25 traffic stream and discards the packets that are out of profile compared to a pre-established traffic profile. When using the Token Bucket procedure, a priority scheduler can be used.

There are several other mechanisms that can be used for policing e.g. WFQ (Weighted Fair Queuing), CBQ (Class Based Queuing) but the fact that they do not
30 provide a strict priority to higher priority classes makes them less useful for real-time applications.

In the article E. P. Ratgeb, "Modelling and Performance Comparison of Policing Mechanisms for ATM Networks", IEEE Journal on selected areas in communications. Vol. 9, No. 3, April 1991, some policing procedures that are "window-based" are
35 discussed and evaluated, including also a "Moving Window" procedure, herein called the Sliding Window procedure.

When using a priority scheduler, there is a risk that traffic of lower priority will suffer from "starvation", i.e. the queues will not be served, as has been mentioned above. This is very critical in networks in which traffic classes of different priorities have real
40 time demands. To guarantee that such starvation will not occur, there must be a

functionality that polices, i.e. meters the traffic flow and discards packets that are out of profile, the traffic flow so that a certain traffic class does not overuse its share of bandwidth. This also protects a traffic class from delaying "itself" too long. Too much traffic in the queue will introduce a too long delay. The IETF proposed the Token Bucket
5 procedure as a policer. The proposed procedure have some properties that are not desirable, especially in the case where the traffic mainly comprises voice traffic. In addition to voice traffic, there are also other real time classes. The undesirable properties comprise that during the time period following immediately after a congestion has been indicated traffic classes of lower priority might suffer from starvation, and that a strict
10 upper limit of the induced delay can not be guaranteed. This results from the fact that the Token Bucket procedure has a reaction time due to the bucket size, which gives a high probability that the bucket is full when congestion is indicated.

Another important aspect is how to measure the aggregated flow when using the Token Bucket procedure. This measurement is required to allow the use of "soft
15 policing". Soft policing means that the policing is only performed when the aggregated flow is above a certain threshold. The Token Bucket procedure has two parameters, the token rate and the bucket size. The token rate is the rate with which the bucket is filled with tokens so it does not tell anything about the flow. The only parameter that tells anything about the flow situation is the present amount of tokens in the bucket, but there
20 is no easy way to use that information to calculate the aggregated flow. This means that some flow meter for the aggregated flow has to be added.

For the Token Bucket procedure to work properly it must be calibrated for the worst case traffic scenario in a network. This might give bad qualities for the other possible traffic cases that can occur in the network.

25 SUMMARY

It is an object of the present invention to provide an efficient policing function for a packet network having values of maximum allowed packet flows.

When using a Sliding Window (SW) based policer function the unwanted properties that the Token Bucket procedure introduces will not trouble the system. A policer function
30 or procedure generally consists of two parts, a meter function or procedure and a dropper function or procedure. The sliding window function is then used as the meter function.

The sliding window procedure can be designed for a network using the fact that there is some information on the way in which the traffic behaves in the network, such as in an IP network. The repetitive characteristic of the codecs and the requirement, that
35 delays should be smaller than one repetition interval, result in that the resource budget is approximately one repetition interval. The bandwidth can therefore be portioned into periods - voice packets are e.g. typically sent every 20 ms - and the resources can be estimated as bytes per period. The use of a sliding window gives a measure of the amount of the resources that are used during this period and packets will be discarded when the
40 resources are over-utilized. In contrast to the Token Bucket procedure, the Sliding

Window procedure sets a fixed upper limit to the amount of traffic that can be transmitted during a certain time. This is an important aspect when the maximum delay is estimated.

By using the sliding window as the meter function in the policer, the special demands that the different traffic classes can have e.g. in IP-networks, especially hard
5 real-time demands, can be fulfilled and at the same time using the bandwidth more efficiently.

With a sliding window based policer it is easy to measure the aggregated flow which implies that packets does not need to be discarded even if a traffic class overutilizes its share of the resources. This is the base for soft policing. Soft policing implies that
10 discarding of packets is only done when the aggregated flow exceeds a certain threshold. This means that the bandwidth of the links can be used more efficiently, what is important for narrow links.

The precision of procedures using a sliding window is maybe the most important feature. The performance of sliding window base procedures is well known - the
15 maximum delay for a traffic class can be calculated - implying that the resources can be used more efficiently.

Generally, when transmitting or discarding arrived data packets in some unit in a network, a sliding window can thus be used. It has a fixed length to meter the flow of data packets and an arrived data packet is transmitted or discarded based on the metered
20 flow. To make this method work more efficiently, it is proposed that the sliding window is divided into a plurality of subwindows of equal length. Counts of transmitted and/or discarded bytes and/or packets are determined both for the sliding window and for each of the subwindows. An arrived data packet is then transmitted or discarded based on the count/s/ both for the current subwindow and for the sliding window. The sliding window
25 can perform its sliding movement by being displaced in steps equal to the length of the subwindows. The movement of the sliding window can be controlled only by the events of arrived data packets. Then, if the current subwindow is defined as that subwindow within which the latest arrived data packet arrive, a new data packet having an arrival time which is not within the current subwindow, the sliding window can thus be moved, the
30 movement step being equal to the length of the subwindows. A new subwindow becomes the current subwindow and has counts set to zero. It replaces the oldest subwindow.

In the case where the arrived data packets belong to different traffic classes having different priorities, a separate sliding window as described can be used for data packets of each traffic class. The sliding window of each class then have a fixed length individual to
35 the class and it also own subwindows, the length and/or number of which can be individual to the class. The count/s/ as described are updated individually for the sliding window and the subwindows of each class. For soft policing, a data packet belonging to a considered class is decided to be allowed to be transmitted or discarded based on a congestion measure such total count/s/ obtained from the count/s/ for the sliding windows
40 of all classes and also preferably from the count/s/ of the current subwindows of all

classes.

A policer module as described herein is thus used for transmitting or discarding arrived data packets and includes metering means executing a sliding window procedure having a fixed window length to meter the flow of data packets. The module further includes decision means, also called a dropper or discarding unit, connected to the metering means to decide whether an arrived data packet is to be transmitted or discarded based on the metered flow. The metering means can then have a plurality of counting means for determining the count/s/ for the sliding window and for the subwindows and memory cells can be provided for storing the counts. Comparing means in the metering means compare the count/s/ of the current subwindow to a maximum value obtained from the count/s/ of the sliding window and an absolute maximum value. The decision means transmits or discards an arrived data packet based on the result of the comparing in the comparing means. Also, for the current subwindow a memory field can be provided for storing the time at which the current subwindow will end and time stepping means can increment this time by a fixed time length when the sliding window is to be moved by one step.

Additional objects and advantages of the invention will be set forth in the description which follows, and in part will be obvious from the description, or may be learned by practice of the invention. The objects and advantages of the invention may be realized and obtained by means of the methods, processes, instrumentalities and combinations particularly pointed out in the appended claims.

BRIEF DESCRIPTION OF THE DRAWINGS

While the novel features of the invention are set forth with particularly in the appended claims, a complete understanding of the invention, both as to organization and content, and of the above and other features thereof may be gained from and the invention will be better appreciated from a consideration of the following detailed description of non-limiting embodiments presented hereinbelow with reference to the accompanying drawings, in which:

- Fig. 1 is a schematic illustrating an IP network,
- Fig. 2 is a schematic illustrating the places where soft and hard or strict policing should be performed in an IP network,
- Fig. 3 is a schematic of a policing module,
- Fig. 4 is a schematic of a sliding window comprising a multitude of subwindows,
- Fig. 5 is a diagram illustrating the effect of different lengths of the sliding window on the momentary bandwidth,
- Fig. 6 is a schematic illustrating some important parameters of a subwindow of a sliding window,
- Fig. 7 is a diagram illustrating the procedural steps executed when receiving a new packet at a policing module in a network comprising one traffic class, and
- Fig. 8 is a diagram similar to that of Fig. 7 for a policing module in a network comprising several traffic classes.

DETAILED DESCRIPTION

The policing and metering functions and devices therefor to be described herein can be used in any kind of packet switched network, e.g. an IP network. In Fig. 1 an IP network 1 is illustrated connecting hosts 3 through routers 5 placed at the border or edge of the network and routers 7 placed inside the network to a gateway 9, which in turn is connected to various services or other networks such as the PSTN (Public Switched Telephone Network) and an IP Core. In the network 1 packets carrying information are transmitted between the hosts and the gateway. The packets can generally belong to different traffic classes requiring different levels of transmission handling in the network, such as the priority, whether packets may be discarded, etc. In the routers 5, 7 congestion of packets can occur, i.e. the routers can receive too many incoming packets and then some received packets must be discarded.

For that reason policer procedures can be included, see the schematic picture of Fig. 2. Here, the hosts 3 transmit packets to edge routers 5 at the border of the network 1 as monitored or supervised by policer procedures as indicated by the rings 13. These procedures, performed at the output side of the transceivers 11 or perhaps preferably at side of the edge routers 5 coupled to the hosts, allow only a maximum flow of packets for each traffic class or generally, the incoming traffic is mapped onto a predefined traffic profile to check if the traffic will get admission to the network or not. In this case where the policing is strictly performed for each traffic class, i.e. where each traffic class is policed individually or independently of the other classes, the policing is called hard or strict policing. Inside the network, at the routers 7 and between the edge routers 5, soft policing is performed as indicated by the rings 15. Soft policing means that the traffic only will be policed if the total band width usage exceeds that of the available physical link. If no congestion exist, more traffic than the configured amount for a considered can be let through. The policing functions, when active, generally operate for every traffic class monitoring all of the different flows all the time. Also, strict policing can in some cases or if so preferred be used inside the network.

Functions for soft and strict policing can be included in the network 1 for instance at the output ports of the routers 7. The policing can then be executed as indicated in Fig. 3. The policing unit or module 17 receives packets from a routing unit 19, in which a routing decision has been made. The flow of received packets is metered, i.e. the statistics of the flow is determined, in the policing unit 17 and based on the metered flow the received packets are either allowed to pass to a scheduler 21, from which the packets are forwarded, or are discarded, as indicated by the trashbin 23. Strict policing can be included in those sides of edge routers 5 which are connected to the hosts or generally to device at edge of the network and then the corresponding policing module include at least the units 17, 21 and 23 of Fig. 3. The discarding of packets can also be made before the packets enter the scheduler 21, in a separate decision or dropping unit connected to the input of the scheduler, as will be described hereinafter.

The policing module 17 can use a sliding window procedure particularly designed for the network, based on information on the traffic behaviour in the network. A sliding window 25 partitioned in N equal subwindows 27 is schematically illustrated in Fig. 4. The sliding window is a time period of length T, during which the packets arrived to the module are evaluated as to their arrival times considering their traffic classes and statistics relating to forwarding and discarding of the packets are determined. The subwindows are thus time periods having the length T/N and for each subwindow also the packets arrived to the module are evaluated and the statistics of the subwindows are determined. The sliding window starts at a time t_s and ends at time $t_e = t_s + T$. The start and end times are changed at regularly repeated times, the interval between these times being a repetition period equal to the length T/N of each subwindow. Thus, for the start times the relation $t_s = n \cdot (T/N) + t_0$ is true, where t_0 is some start time and n is an integer value incremented every (T/N) :th unit of time. Hence, the sliding window slides in time or more exactly makes small jumps or hops having a length of a subwindow.

A new repetition period thus starts with every new subwindow, which thus is the current subwindow and has the order number N. The amount of resources, i.e. the maximum number of bytes or the maximum number of packets which are allowed to be transmitted, to be used for the N:th subwindow is the amount corresponding to the maximum resources to be used for the whole window subtracted by the resources used in the other, N-1 previous subwindows. When the resources have been spent in the Nth subwindow, the new packets that arrive within this window are discarded, i.e. the new packets are discarded until a new period or subwindow starts which will then be the Nth subwindow. This guarantees that no more than a certain amount of bytes or packets is spent during the time T of the whole sliding window.

In the diagram of Fig. 5 an example of a varying momentary bandwidth of incoming traffic is shown as a function of time, the momentary bandwidth being essentially the flow of bytes or packets, i.e. the number of bytes or packets received per unit of time. It appears that the method works well in the case where there is one traffic peak and one traffic valley during one sliding window, see the arrow 29 indicating a suitable length of the total sliding window. For a too short sliding window, see the arrow 31, which covers just one peak or one valley, too many packets are discarded when a traffic peak occurs and the link is then also poorly used when a valley occurs.

For each subwindow the determined statistics are stored in corresponding memory fields as will be described hereinafter. The statistics comprise several parameters describing the packet flow during the subwindow, see Fig. 6. The most important information to be stored comprises the numbers of bytes and the numbers of packets that are allowed to be transmitted and are discarded in each subwindow. The information can also comprise a parameter indicating the time when the present subwindow is to be/was replaced, the window limit. Preferably, the updating of the sliding window, i.e. the change of subwindows, is not driven by an interrupt but instead is event driven when a

packet arrives, the arrival time of the packets must be compared to the window-limit to check whether the statistics for the current window or the next one has to be updated. However, the updating can also, obviously, be hardware or software driven in other ways such as by hardware or software interrupts or monitored in a parallel software process.

5 Hence, when a packet arrives, the arrival time of the packet is determined and use to decide the subwindow to which the packet belongs. If the arrival time is smaller than the window limit, the statistics of the present subwindow is updated. If the arrival time of the packet is greater than the window limit, the statistics for the total sliding window must be updated. Then the statistics for the oldest subwindow in the window are subtracted from
 10 the statistics of the total sliding window, and the statistics of the latest subwindow, i.e. the subwindow which is being terminated, is added to the statistics of the total sliding window. Sometimes the time elapsed between two consecutive packets can be longer than one subwindow. Then the updating continues iteratively to the next subwindows until the window limit of a subwindow is greater than the arrival time of the packet.

15 The handling and procedural steps performed in the policer module also appear from Fig. 7 for the case of only one traffic class. In this case, according to the discussion above, the policer module executes only hard or strict policing. In a timestamping block or unit 41 packets are received and the time is determined of the arrival of each packet. Also the number of bytes of the packet can be determined for a later updating of the
 20 statistics of the current subwindow. In a decision block or unit 43, also called a discarding or dropper unit, it is decided whether a packet will be transmitted or forwarded, see the arrow 45, or be discarded, see the arrow 47. In the case where the packet is to be forwarded, it is stored in buffer of a scheduler, see Fig. 3, from which the buffered packets are transmitted in some chosen order. The metering functions of the policer
 25 module are performed by a metering unit 49.

In a memory 51 of the policing module memory cells C_1, \dots, C_N are provided for storing the statistics of the subwindows of the sliding window, the memory cell C_N storing statistics of the current subwindow. Each memory cell includes fields, not shown, for the numbers or counts described above, i.e. counts for forwarded bytes, counts for forwarded
 30 packets, counts for discarded bytes, and counts for discarded packets. In at least the current memory cell C_N there is a field for storing the time limit, i.e. a value indicating the time when the subwindows will terminate. The memory cells C_1, \dots, C_N are preferably arranged cyclically, as a cyclical memory, the memory cell C_N of the current window being pointed to by a pointer stored in a pointer register PR. The memory 51 also
 35 has a memory cell C_{SW} for the statistics of the whole sliding window. This memory cell holds fields, not shown, for the same kind of counts as described above.

At the instant when a new packet is timestamped by the block 41, a sequence of procedural steps are performed. Thus, it is first determined in a block 61 whether the packet has arrived within the current subwindow. This is made by comparing the arrival
 40 time of the packet to the limit of the current subwindow as stored in a field of the

memory cell C_N assigned to the current subwindow. If the received packet is determined to be within the current subwindow, it must be decided whether already the allowed number of packets and/or bytes has been transmitted during this subwindow and thus a block 63 is executed. In this block the different counts of the current subwindow are compared according to some algorithm to maximum values. The maximum values are in turn derived from the statistics of the total sliding window as stored in the memory cell C_{SW} and absolute maximum values. In a simple case/According to a simple algorithm only the count of forwarded bytes of packets for the current subwindow is compared to the difference of the absolute maximum count for forwarded packets and the count of forwarded packets of the total sliding window. If it is decided in the block 63 that the packet can be forwarded, a block 65 is executed in which the statistics of the current subwindow are updated. This means that the fields of the counts of forwarded packets and forwarded bytes in the cell C_N are incremented by one and by the number of bytes of the packet respectively. A signal is then sent from the next block 67 to the decision unit 43 informing it that the packet is to be forwarded. After the block 67 the arrival of a new packet is awaited in the timestamping unit 41.

If it is decided in the block 63 that the packet cannot be forwarded but has to be discarded, a block 69 is executed for updating the statistics of the current subwindow. Then the fields of the counts of discarded cells and discarded bytes in the cell C_N are incremented by one and by the number of bytes of the packet respectively. In the next block 71 a signal is then sent to the decision unit 43 informing it that the packet is to be discarded. After the block 71 the arrival of a new packet is awaited in the timestamping unit 41.

In the case where it is determined in the block 61 that the new packet has an arrival time not within the current subwindow, the sliding window must be allowed to "slide" or be displaced forwards by the length of at least one subwindow. Then a block 73 is executed in which first the statistics of the total sliding window are modified. This is made as described above by subtracting the counts for the oldest subwindow as stored in the fields of the memory cell C_1 from the respective counts of the total sliding window and then adding the counts of the current window to the result of the subtraction. The pointer value is changed to the next position in the cyclical sequential order of the memory cells C_1, \dots, C_N , i.e. to point to the memory cell of the oldest subwindow. The statistics of this subwindow are reset to zero and a new time limit is calculated and stored in the field provided therefor. Then the block 61 is again executed to determine whether the packet is within the now selected, current subwindow.

It is obvious that the procedural steps of Fig. 7 can be modified in different ways and arranged with different logic decisions. Some of the steps can e.g. be performed in an inverted order or simultaneously.

For the case where several different traffic classes are used in the network, such as for soft policing, each traffic class can be policed by its own sliding window where the

length of the sliding window is set to match either the repetitive characteristics of the traffic or by the maximum delay requirements of the packets of the class. Congestion is indicated when a weighted sum of the traffic sent during the different sliding windows exceeds a certain threshold. After detecting a congestion state, either every class can be
 5 policed separately using for each class independently the method described with reference to Fig. 7, or a distribution algorithm can be used according to which a certain amount of resources will be available for the different sliding windows as long as the congestion situation persists. The algorithm can be designed to distribute in a flexible manner the amount of resources to the different SWs.

10 The procedural steps to be executed in this case appear from Fig. 8. For each traffic class P_1, P_2, \dots memory cells $C_{1,Pn}, C_{2,Pn}, \dots, C_{N,Pn}$, etc., $n = 1, 2, \dots$, are provided, and it is assumed that the sliding windows of the traffic classes all have the same number N of subwindows. The subwindows then can have different lengths for the classes. However, it is also possible to have different numbers of subwindows but this case is not
 15 discussed with reference to Fig. 8. A pointer to the current subwindow for each class is provided. The updating of the different sliding windows can then be made for the shortest one of the subwindows. Preferably, every subwindow for every class should start at the time when one of the shortest windows starts. This means i.a. that in this case the lengths of the subwindows are multiples of the length of the shortest subwindows.

20 The basic units involved in the handling are only modified by the timestamping unit 41 also determining the traffic classes of the received packets. When the timestamping unit 41 detects the arrival of a new packet, a signal is transmitted to start the logical steps of the procedure. Then a first decision block 61' is executed in which it is asked whether the time of arrival of the packet is within the smallest of the current subwindows. If it is
 25 true a block 81 is executed in which it is decided whether the module is in congestion mode, as will be described below. If the policing module is decided not to be in a congestion mode, a block 83 is executed in which basically the same operations are performed as in the block 65 of Fig. 7 but for the fields of the memory cell of the current subwindow for the class to which the packet belongs. Then a block 85 similar to the block
 30 67 is executed whereafter the arrival of new packet is again awaited in the unit 41.

If it was found in the block 81 that the module is in a congestion mode a block 87 is executed in which it is determined whether there still exist resources for the traffic class to which the packet belongs. This means that substantially the same comparison is made as in block 63 of Fig. 7, i.e. evaluating the counts of the current subwindow of the class
 35 of the packet in relation to the counts for the total sliding window of the class and to maximum allowed counts for the class. If it is determined that resources are still available for the class, blocks 89 and 91 are executed in which the same operations are performed as in blocks 83 and 85 respectively described above. Then the arrival of a new packet is again awaited in the unit 41.

40 If it is decided in the block 87 that no resources are left for the class, blocks 93 and

95 are executed in which the same operations are performed as in blocks 69 and 71 of Fig. 7 but the updating of statistics is made for the subwindow of the traffic class of the packet. Then the arrival of a new packet is again awaited in the unit 41.

If it is decided in the block 61' that the arrival time of the packet is not within the
5 smallest of the current subwindows a new decision is made in a block 61" in which it is determined whether the arrival time of the packet is within the current subwindow of the traffic class having the first order number $n = 1$. If the arrival time is within this subwindow, a block 97 is executed in which the order number n of the traffic class is incremented by one. If there is such a class, i.e. if the new $n \leq N_p$, the block 61" is
10 again performed for the new n , i.e. it is decided whether the arrival time is within the subwindow of class n . If there is no such another class, after incrementing n in block 97, i.e. the new $n > N_p$, a value of the aggregated flow is calculated in a block 99, e.g. as a weighted sum of the counts for all sliding windows. The counts used in this weighted sum should be counts indicating in the best possible way the current state for the classes, and
15 can be obtained from the statistics of the current total sliding windows to which are added the statistics of the current subwindows of the respective class. The value of the aggregated flow is used for determination of a possible congestion state in the block 81 described above. In the case where the value indicates a congestion state, e.g. if the aggregated flow value exceeds some threshold value, a block 101 is executed in which the
20 resources for the traffic classes are calculated using a policing algorithm. This algorithm can distribute the remaining resources in a static, predetermined way used fixed shares for each class, whereafter strict policing for each class is performed, until the congestion state disappears. After the block 101 and in the case where no congestion state is indicated the block 81 is performed as described above.

25 It can be determined in the block 61" that the arrival time is not within the subwindow of class n and then a block 103 is executed, in which the statistics and the pointer are modified as described for the block 73 of Fig. 7 but only for the memory cells belonging to the traffic class n . Then the block 61" is executed again as described above.

While specific embodiments of the invention have been illustrated and described
30 herein, it is realized that numerous additional advantages, modifications and changes will readily occur to those skilled in the art. Therefore, the invention in its broader aspects is not limited to the specific details, representative devices and illustrated examples shown and described herein. Accordingly, various modifications may be made without departing from the spirit or scope of the general inventive concept as defined by the appended
35 claims and their equivalents. It is therefore to be understood that the appended claims are intended to cover all such modifications and changes as fall within a true spirit and scope of the invention.

CLAIMS

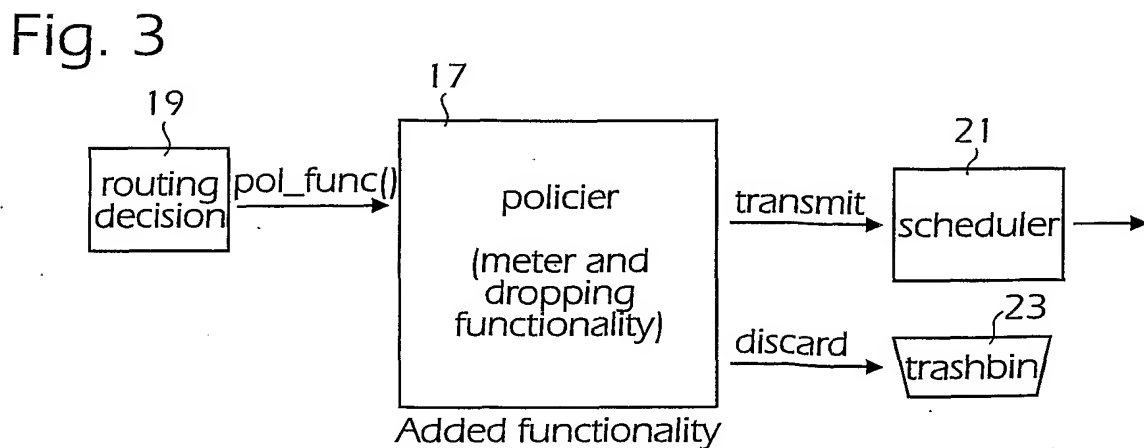

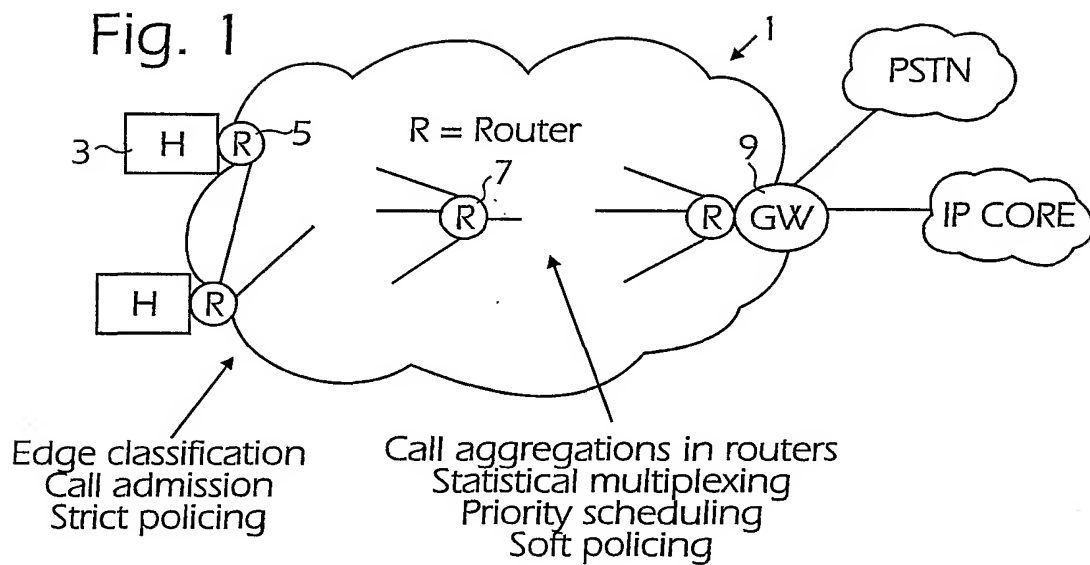
1. A method of transmitting arrived data packets, the method including the use of a sliding window having a fixed length to meter the flow of data packets, an arrived data packet being transmitted or discarded based on the metered flow, **characterized in** that
5 the sliding window is divided into a plurality of subwindows of equal length including a current subwindow and an oldest subwindow, that counts of transmitted and/or discarded bytes and/or packets are determined for the sliding window and the subwindows, and that a data packet is transmitted or discarded based on the count/s/ both for the current subwindow and for the sliding window.
- 10 2. A method according to claim 1, **characterized in** that the sliding window moves by steps equal to the length of the subwindows.
3. A method according to claim 1, **characterized in** that for a data packet having an arrival time not within the current subwindow, the sliding window is moved by a step equal to the length of the subwindows, a new subwindow becoming the current
15 subwindow having counts set to zero.
4. A method according to claim 3, **characterized in** that when the sliding window is moved, the count/s/ of the sliding window is/are modified by subtracting count/s/ of the oldest window and adding count/s/ of the formerly current subwindow.
5. A method according to claim 1, **characterized in** that for an arrived data packet
20 the count of forwarded bytes or packets for the current subwindow is compared to the difference of an allowed maximum number of forwarded bytes or packets respectively and the count of forwarded bytes or packets respectively for the sliding window.
6. A method according to claim 1, in the case where the arrived data packets belong to different traffic classes having different priorities, **characterized in** that a separate
25 sliding window is used for data packets of each traffic class, each separate sliding window having a fixed length individual to the class and own subwindows and the count/s/ updated individually for the sliding window and the subwindows of each class, a data packet belonging to a considered class being transmitted based on a total count/s/ obtained from the count/s/ for the sliding windows of all classes.
- 30 7. A method according to claim 6, **characterized in** that the total count/s/ are obtained also from the count/s/ of the current subwindows of all classes.
8. A method according to claim 6, **characterized in** that for an arrived data packet a value of aggregated flow is calculated as a weighted sum of the count/s/ for the sliding windows of all classes and for the current subwindows of all classes.
- 35 9. A policer module for transmitting or discarding arrived data packets, the module including metering means executing a sliding window procedure having a fixed window length to meter the flow of data packets, and further including decision means connected to the metering means to decide whether an arrived data packet is to be transmitted or discarded based on the metered flow, **characterized in** that the metering means comprise
40 counting means for determining counts of transmitted and/or discarded bytes and/or




packets of the sliding window and for a plurality of subwindows of the sliding window, the subwindows having equal length and including a current subwindow and an oldest subwindow, and comparing means for comparing the count/s/ of the current subwindow to a maximum value obtained from the count/s/ of the sliding window and an absolute maximum value, the decision means arranged to transmit or discard an arrived data packet based on the result of the comparing.


10. A module according to claim 9, in the case where the arrived data packets belong to different traffic classes having different priorities, **characterized in** that the metering means are arranged to use a separate sliding window for data packets of each traffic class and having a length individual to the class, each separate sliding window having own subwindows and that the counting means are arranged individually for the sliding window and the subwindows of each class, the comparing means arranged, in comparing, to use a value of the total transmitted flow obtained from the count/s/ for the sliding windows of all classes and compare this value to a set maximum value.

11. A module according to claim 10, **characterized in** that the value of the total transmitted flow is obtained also from the count/s/ of the current subwindows of all classes.

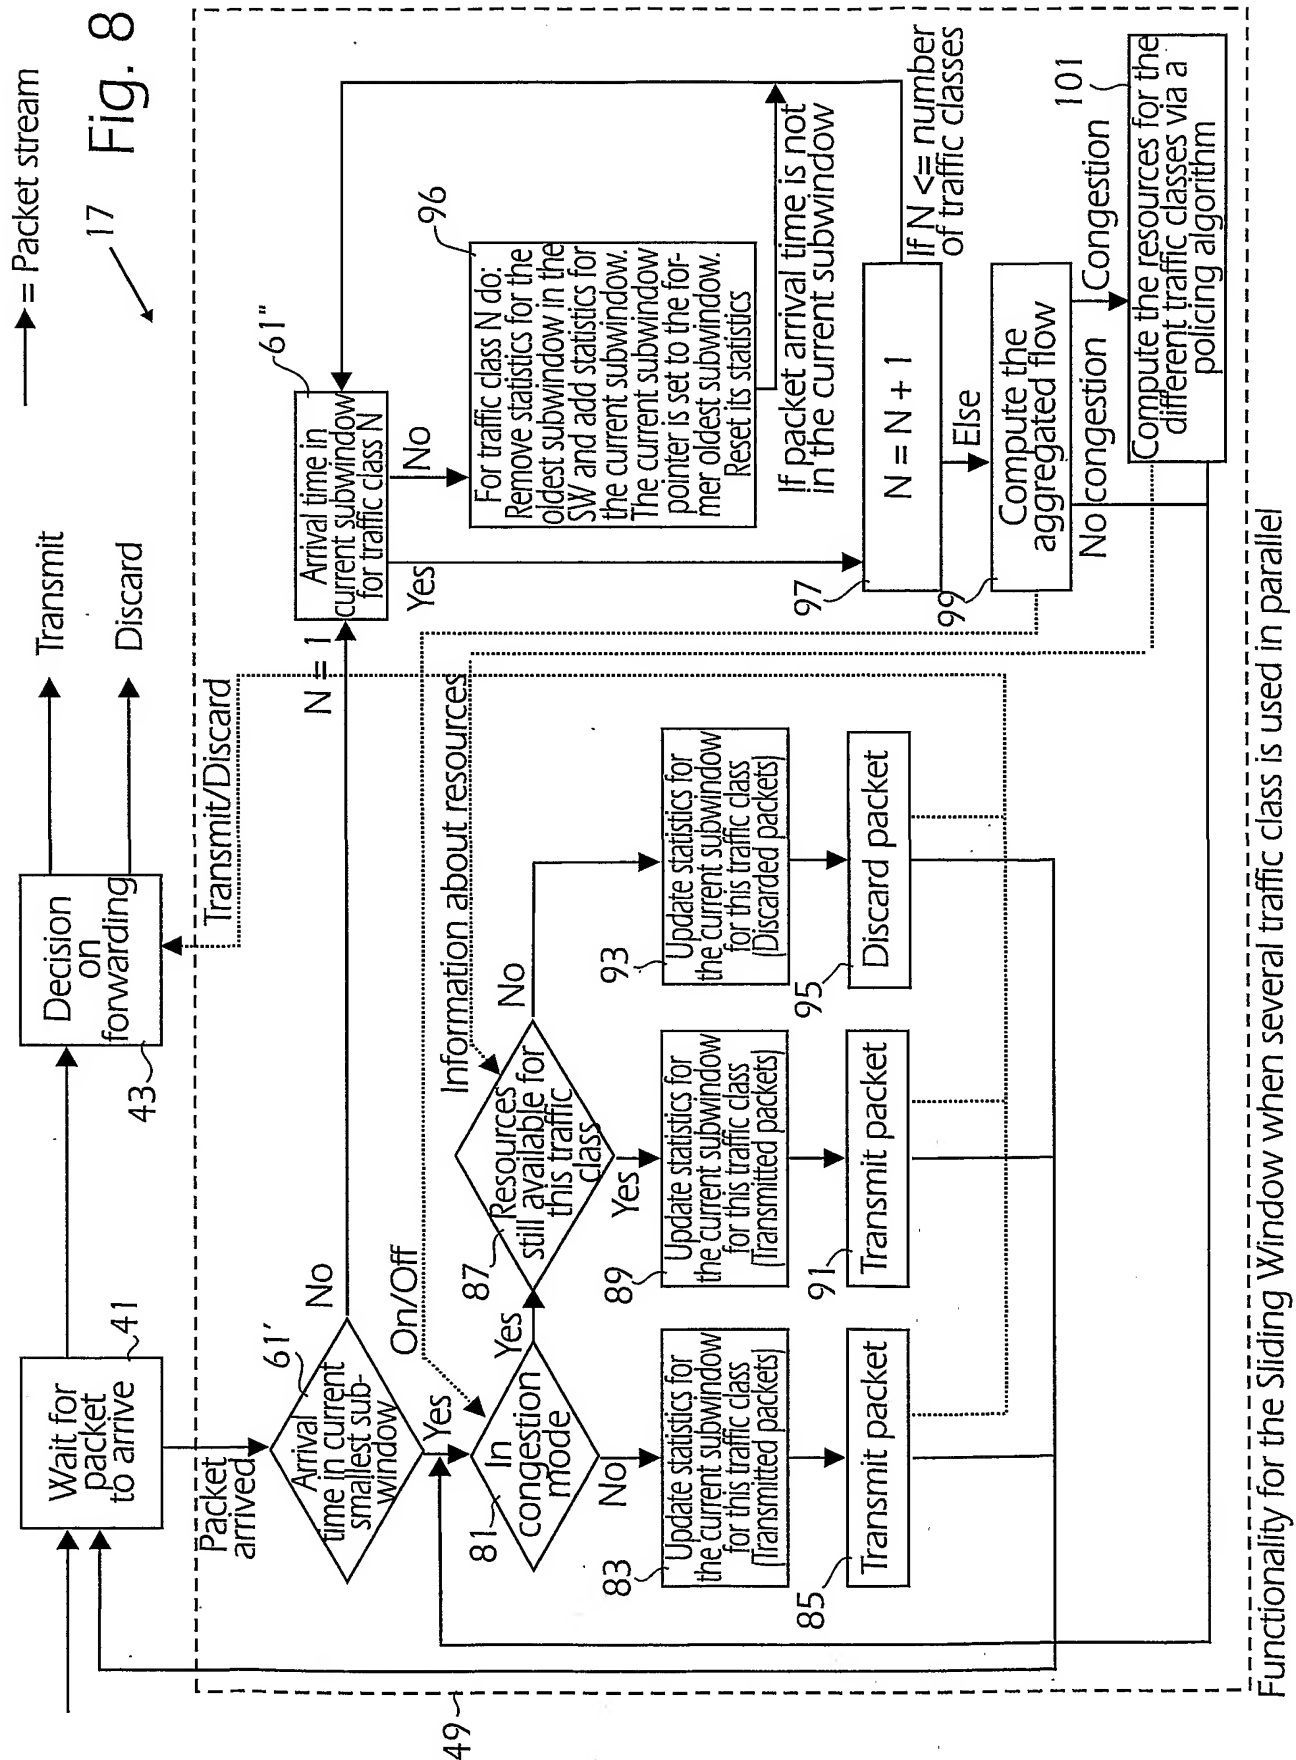
12. A data packet network comprises a policer module for transmitting or discarding arrived data packets, the module including metering means executing a sliding window procedure having a fixed window length to meter the flow of data packets, and further including decision means connected to the metering means to decide whether an arrived data packet is to be transmitted or discarded based on the metered flow, **characterized in** that the metering means comprise counting means for determining counts of transmitted and/or discarded bytes and/or packets of the sliding window and for a plurality of subwindows of the sliding window, the subwindows having equal length and including a current subwindow and an oldest subwindow, and comparing means for comparing the count/s/ of the current subwindow to a maximum value obtained from the count/s/ of the sliding window and an absolute maximum value, the decision means arranged to transmit or discard an arrived data packet based on the result of the comparing.







Functionality for the Sliding Window when only one traffic class is used



INTERNATIONAL SEARCH REPORT

International application No.

PCT/SE 01/02583

A. CLASSIFICATION OF SUBJECT MATTER

IPC7: H04L 12/56, H04Q 11/04

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC7: H04L, H04Q

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

SE,DK,FI,NO classes as above

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-INTERNAL, WPI DATA, PAJ, INSPEC, COMPENDEX

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	CA 2249152 A1 (NORTEL NETWORKS CORP), 30 March 2000 (30.03.00) --	1-12
A	WO 9636150 A1 (NOKIA TELECOMMUNICATIONS OY), 14 November 1996 (14.11.96) --	1-12
A	WO 0056023 A1 (TELEFONAKTIEBOLAGET LM ERICSSON (PUBL)), 21 Sept 2000 (21.09.00) --	1-12
A	US 5930233 A (KANERVA, M. ET AL), 27 July 1999 (27.07.99) -- -----	1-12



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"I" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance: the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance: the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

Date of mailing of the international search report

25 February 2002

Name and mailing address of the ISA/

Swedish Patent Office

Box 5055, S-102 42 STOCKHOLM

Facsimile No. +46 8 666 02 86

Authorized officer

Marianne Norrgren/LR

Telephone No. +46 8 782 25 00

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/SE 01/02583

Patent document cited in search report			Publication date	Patent family member(s)		Publication date
CA	2249152	A1	30/03/00	NONE		
WO	9636150	A1	14/11/96	AU	707904 B	22/07/99
				AU	5650296 A	29/11/96
				CA	2220426 A	14/11/96
				CN	1183870 A	03/06/98
				EP	0826280 A	04/03/98
				FI	98023 B,C	13/12/96
				FI	952255 D	00/00/00
				JP	11505084 T	11/05/99
WO	0056023	A1	21/09/00	AU	3852700 A	04/10/00
				SE	0102982 A	10/09/01
US	5930233	A	27/07/99	AU	708324 B	05/08/99
				AU	5650696 A	29/11/96
				CA	2193379 A	14/11/96
				EP	0788702 A	13/08/97
				FI	98174 B,C	15/01/97
				FI	952256 A	10/11/96
				JP	10503075 T	17/03/98
				WO	9636154 A	14/11/96